

REMARKS

Claims 1, 4, 6-13, 15-18 and 21-26 are pending in the patent application.

The Examiner has objected to the Specification, has rejected Claims 1, 10, 12 and 21-26 under 35 USC 112 and has finally rejected Claims 1, 4, 6-13, 15-18 and 21-26 under 35 USC 103(a) as being unpatentable over Yang in view of Chen and further in view of Miller.

In objecting to the Specification, the Examiner has stated that the specification fails to provide proper antecedent basis for the term *existing* workload units. The Applicants have amended the claims to instead recite "current" workload units, as is clearly supported by the Specification (see: page 15, line 25 and page 18, lines 12-13). Applicants reiterate that the present invention is directed to assessing and reassigning workload units in a computer system which is up and running, as opposed to cited prior art which is directed to designing a computer system to handle expected future workload. Applicants believe that the use of the term "current" in the claims serves to clearly differentiate the invention over the prior work done in computer system modeling and design. Withdrawal of the 112 rejection is accordingly requested.

In the Response to Arguments section, the Examiner cites several definitions of the term "workload" in an attempt to support the Examiner's conclusion that "a workload unit is a unit of work to be performed in the future". Applicants note that the Examiner has not provided any source of the definitions and has clearly edited the definitions by abbreviating them. As such, the definitions cannot support the Examiner's conclusion. Moreover, Applicants contend that the Examiner's conclusion is erroneous. Applicants' attorney found several definitions of the term "workload" that contradict the Examiner's conclusion and support a conclusion that "workload" can refer to a current or a projected amount of work. Answers.com defines "workload" as

1. The amount of work assigned to or expected from a worker in a specified time period.
2. The amount of work that a machine produces or can produce in a specified time period.

The Federal Circuit has stated that the meaning of a particular claim term may be defined by implication, that is, according to the usage of the term in the context in the specification. See, e.g., Phillips v. A. WH. Corp., 415 F.3d 1303, 75 USPQ2d 1321 (Fed. Cir. 2005); and Vitronics Corp. v. Conceptronic Inc., 90 F.3d 1576, 1583, 39 USPQ2d 1573, YOR920000461US1

1577 (Fed. Cir. 1996). See generally MPEP 21 11.01. Moreover, the Patent and Trademark Office (PTO) determines the scope of claims in patent applications not solely on the basis of the claim language, but upon giving claims their broadest reasonable construction 'in light of the specification as it would be interpreted by one of ordinary skill in the art." MPEP 2111 (citing In re Am. Acad. of Sci. Tech. Ctr., 367 F.3d 1359, 1364, 70 USPQ2d 1827 (Fed. Cir. 2004)). Since it is clear from the present Specification that "workload" refers to work that is currently assigned to a computer system (see: again page 15, line 25 and page 18, lines 12-13), that definition of "workload" must prevail. Applicants request entry of the amendments and withdrawal of the objection to the Specification.

In response to the 112 rejections, Applicants have amended the language of Claim 25 to appropriately recite "the key service date". In addition, the language of the independent claims has been amended to address the Examiner's concern with the term "existing workload". Support for the amendments is found in the original Specification (see: e.g., Fig. 10 and accompanying text on pages 20-21).

In the **Response to Arguments** section, the Examiner states that "the specification does not detail how a process
YOR920000461US1

being processed is moved or reassigned, the process must be a queue or the like", Applicants disagree and respectfully direct the Examiner's attention to Fig. 10 and the accompanying description on pages 20-21.

In response to the rejections under 35 USC 103(a) as unpatentable over Yang in view of Chen and Miller, Applicants respectfully traverse the rejections. Applicants assert that the Chen patent does not qualify as prior art against the present application. All rejections under 35 USC 103 logically rely on the relationship of the applied prior art to the claimed invention. The pending rejection uses various references in view of Chen to deny patentability of the present invention under 35 USC 103, because of an assumption that Chen would qualify as prior art under 35 USC 102. However, this assumption is incorrect. An inspection of Chen reveals that the Chen patent is currently assigned to International Business Machines Corporation (IBM). The subject matter of Chen was owned by IBM and subject to an obligation of assignment to IBM at the time that the present invention was made. In addition, at the same time, the present invention was owned by IBM and subject to an obligation of assignment to IBM. Hence, the present invention and the subject matter of Chen were owned by a common assignee (i.e., IBM) at the time that the present invention was made.

The American Inventors Protection Act (AIPA) of 1999 changed 35 USC 103(c), which applies to any patent application filed on or after the date of enactment (11/29/99). 35 USC 103(c) states:

(c) Subject matter developed by another person, which qualifies as prior art only under one or more of subsections (e), (f), and (g) of section 102 of this title, shall not preclude patentability under this section where the subject matter and the claimed invention were, at the time the invention was made, owned by the same person or subject to an obligation of assignment to the same person.

In other words, prior art that would otherwise qualify under 35 USC 102(e), which requires that "the invention was described in (1) an application for patent... or (2) a patent granted on an application for patent by another filed in the United States before the invention by the applicant", cannot be used to deny patentability if the prior art and the claimed invention were commonly owned or subject to an assignment to a common assignee at the time of the invention. 35 USC 103(c) is applicable to the present invention because the present application was filed on 10/19/2000. Accordingly, Chen is disqualified as prior art against the present invention. Since the grounds of the

YOR920000461US1 -20-

pending rejection can no longer be based on the patent of Chen, the pending ground of rejection must be withdrawn.

Applicants will, nonetheless, address the Examiner's rejections and interpretations of the applied art. The present invention is directed to an apparatus, program storage device, and a method for evaluating workload across a processing environment having a plurality of computer systems each having a plurality of assigned current workload units, wherein the method comprises the steps of assigning a plurality of impact values, one impact value for each current workload unit assigned for each of the plurality of computing systems, wherein the assigning of each impact value comprises determining a change in system expiration date should the workload unit be removed from the system; and assessing the workload based on the impact values. An impact value is assigned for each workload unit, wherein a workload unit is expressly defined for the application as "a subset of the workload", the workload being "the set of identifiable tasks that execute in the processing system" (see: page 8, line 19-page 9, line 3). For each subset of the workload, an impact value is assigned, representing the change in system expiration date that would occur if the workload unit was removed from the system. The term "expiration date" is the date when the server workload is expected to exceed its capacity because of growth in

YOR920000461US1 -21-

workload. The expiration date may be calculated using life expectancy, capacity space, or other method, as detailed in the Specification on page 11, lines 1-18.

The Yang patent is directed to a method and apparatus for modeling or profiling a **future** system based on **anticipated** workload in order to **design** the system. A model is generated using "a set of generic system activities" and performance estimates, after which hardware parameters are determined to **design** a system to handle the **expected** activities. Yang uses computer activity elements ("CAE") "to capture the fundamental activities being performed by [a] respective software application" (see: Col. 4, line 1-8). A set of CAE is calculated for each transaction (see: Col. 26, lines 18-19). A "user focused workload" ("UFW") is collected "using forms or templates ...regarding the application that is available to the user" (Col. 5, lines 50-53). The UFW is translated into a computer activity workload ("CAW") representing the workload data structure that mathematically describes the workload in terms of CAE elements (Col. 5, lines 58-62). Yang **models expected** application workload for a user and then recommends system requirements for creating a system for that specific user for the user's **expected usage** of the system. As is clearly taught by Yang, the hardware requirements are determined

based on **modeled workload**. Step 112 in Fig. 1 expressly shows that the final step of Yang's process is to determine hardware parameters. Similarly, step 410 of Fig. 4 shows a final step to "profile hardware system" and step 510 in Fig. 5 shows the final step to "determine possible hardware configurations". **Yang in not analyzing an existing processing environment with computer systems having assigned current workloads and trying to determine how to re-distribute workload among them.** Rather, Yang is modeling anticipated workload and designing a system to accommodate that workload. Yang does not determine the impact value of current workload units in order to juggle those workload units among currently working computer systems. Yang determines anticipated CAEs and then recommends an appropriate amount of hardware resources to accommodate the number of CAEs in the projected workload.

Applicants respectfully assert that the Yang patent neither teaches nor suggests the invention as claimed. With specific reference to the language of the independent claims 1, 10, 12, 21, 23, and 25, Yang does not teach or suggest evaluating current workload across an actual processing environment having a plurality of computer systems each having a plurality of assigned current workload units.

Rather, Yang models anticipated requirements for projected applications under user-estimated usage conditions.

The Examiner has, in the **Response to Arguments** section, stated that Yang teaches sizing a system "with respect to the anticipated or actual 'workload' that is to be performed by a software application" citing Yang at Col. 3, lines 59-61. Further citations from Yang refer to collecting "a running instance of a workload" and "a set of actions that is performed by the software application". Applicants reiterate that Yang is planning a future system based on projected workload. Yang may use "workload definition information collected from other sources" (Col. 6, lines 9-10) such as the "running instance of a workload" which is running on a different system...not the one being designed. Clearly Yang is teaching **projecting** a workload, whether based on actual workload from another source or anticipated workload from pure conjecture. When Yang takes into account events which are "added to the workload, or even preempt the current workload", as cited by the Examiner, Yang is predicting behaviors as well as likely changes and interruptions that can be expected to occur to the workload. Yang is using standard modeling techniques to design a system and is not teaching or suggesting assessing a current workload on a currently running system.

In rejecting the steps and means for assigning a plurality of impact values to assigned workload units (Claims 1, 10, 12 and 22, 24 and 26), the Examiner generally cites the teachings from Col. 5, line 1 to Col. 6, line 19. Applicants respectfully assert that the cited passage does not teach or suggest the assigning of impact values. Yang generates a workload data structure to mathematically describe workload estimates by assigning CAEs in a processing "vacuum", such that an existing processing environment and existing computer system's workload is not taken into account. In contrast, and as specifically taught by the present Specification at page 5, line 7-9, the measure/impact value of a workload unit "is calculated in relation to a specific processing system and the other workload assigned to the system". Where the current claims expressly recite the assignment of impact values relative to the existing workload and system expiration date, Yang uses theoretical values for defining CAEs and the amount of resources required for CAEs (see: the cited passage from Col. 5, line 1 to Col. 6, line 19). While Yang may use prior/historical averages of resource requirements based on data from other systems (Col. 6, lines 9-12), Yang is neither teaching nor suggesting determining impact values for an existing system having an assigned workload.

In rejecting the independent claim features of the steps and means for assessing the impact of moving a workload unit from a donor computer system to a recipient computer system based on impact values (Claims 1, 10 and 12 and Claims 21-26), the Examiner has again cited the passage from Col. 33, lines 30-62. What Yang teaches in the cited passage is that different possible hardware configurations are considered for a modeled workload. Yang looks at the resources of a hardware configuration and determines whether the resources are sufficient to handle the modeled CAW. If not, Yang looks at a different projected hardware configuration to handle the CAW. If more than one hardware configuration is suitable, Yang may further look at other "desirability" factors (e.g., cost) in selecting a hardware configuration. Yang neither teaches nor suggests assessing the impact of moving some of the CAEs of the CAW to a different hardware configuration.

Yang does not teach or suggest that a system has an expiration date, since Yang does not have a predefined, existing processing environment having a plurality of computer systems each having a plurality of assigned workload units. Rather, Yang models a larger system to accommodate larger workloads, so the concept of system expiration date is simply not applicable. Applicants again note that the Examiner has not cited any teachings in Yang YOR920000461US1

against the "system expiration" language. Rather, the Examiner states "the use and advantages for using such a system is well known to one skilled in the art at the time the invention was made as evidenced by the teachings of Chen".

Applicants reiterate that Chen is not available as a reference against the present application. Applicants further reiterate that the rejections cannot be sustained based on a combination of teachings from Yang and Chen. The Chen patent is directed to a system having parallel database locations for storing data. When a new node (i.e., storage location) is added to the system or when an existing node has excess capacity (see: Abstract), data is moved from one storage node to the new or existing available node. The Examiner has cited Col. 4, lines 56-61 of Chen as providing the teachings missing from the Yang patent, namely "a processing environment each computer system having a plurality of assigned workload units and altering the workload in the processing environment to change expiration dates of at least two of said plurality of computer systems". The cited passage from Col. 4 of Chen merely states that multiple computer systems are coupled to a network, but does not teach or suggest workload units, altering workload, or changing expiration dates of computer systems. The Examiner additionally cites Col. 6, lines 3-32 YOR920000461US1

of Chen. The cited Col. 6 passage describes the data redistribution process of Fig. 6 but states that "selecting the buckets to be redistributed is outside the scope of this present invention" citing a prior article and a patent application. The passage does not teach or suggest computing system workload units, altering workload among a plurality of computing systems, or changing expiration dates of computer systems. Mention is made of selecting buckets of stored data for redistribution "to minimize the impact to the performance (*sic*) of the PDB system" and that value "is measured by determining the amount of workload in a bucket and the cost of moving the bucket from node to node...for data distribution". Workload in a bucket of stored data is not defined by Chen; however, it clearly does not refer to processing workload, since Chen is moving buckets of stored data and not actual processing workload. Applicants conclude that the Chen teachings of moving data to a different data storage device is not the same as or suggestive of moving existing processing workload from one computing system to another based on computing system expiration date.

Applicants also contend that if one skilled in the art were to modify Yang with Chen, the result would be a Yang modeling system which uses Chen's approach to predict how much storage will be necessary for storing data in the

YOR920000461US1 -28-

system which is being modeled. One would not arrive at a system which selects data to be moved from one data storage location to another, since Yang is planning a system and would not yet have any data to store. Clearly there would not be excess capacity in data storage locations since Yang would not yet have assigned an amount of storage capacity.

In acknowledging that Yang and Chen fail to teach assigning an impact number representing the number of days that the expiration data of the computer system would be changed with all other workload units remaining the same, the Examiner has newly cited the Miller patent. Miller teaches modifying tasks, themselves (Col. 3, lines 29-37), to require less computing capability. Miller does not teach or suggest assigning impact values to workload units as they relate to system expiration date and then assessing the effect of moving workload units from one computer system to another. Rather, Miller expressly teaches modifying the workload units, themselves, (see: Fig. 2, step 52) so that the workload units require less processing time or fewer resources. The cited passage from Col. 9, lines 3-68 teach the initial Miller assessment of task requirements and resource availability, but do not teach or suggest assigning impact values related to system expiration date. Further, since Miller teaches away from moving workload among resources and instead teaches modifying tasks at assigned

YOR920000461US1 -29-

resources, it cannot be concluded that Miller provides the teachings which are missing from the combination of Yang and Chen.

For a determination of obviousness, the prior art must teach or suggest all of the claim limitations. "All words in a claim must be considered in judging the patentability of that claim against the prior art" (In re Wilson, 424 F. 2d 1382, 1385, 165 U.S.P.Q. 494, 496 (C.C.P.A. 1970)). Since the cited Yang and Chen references fail to teach each and every one of the claim limitations, a *prima facie* case of obviousness has not been established by the Examiner.

With specific reference to the language of the dependent claims, Applicants note that the Examiner has repeated the previous rejections and citations. Applicants, therefore, maintain the contention that the combination of references does not obviate the claims.

With regard to Claims 4 and 13, the Examiner has again cited the passage, and code detailed therein, from Col. 26, line 45-Col. 27, line 5. Yang detailed calculating anticipated reads and writes for one **projected** transaction to **model** the workload for the transaction. Such is not the same as or suggestive of sorting actual current workload units based on assigned impact values.

With regard to the language of Claims 5 and 14, which has been added to the independent claims by this amendment, the Examiner again cited the passage from Col. 25, lines 13-20 which teaches characterizing one transaction in terms of CAE. There is no mention in the passage of altering current workload of one or more system. Applicants reiterate that since Yang is "operating" in a theoretical environment, Yang neither teaches nor suggests actual expiration dates of actual systems or altering actual current workloads for one or more systems.

With respect to Claims 6 and 15, and Claims 7-9 and 15-18 which depend respectively therefrom, Applicants argue that the cited passage from Col. 33, lines 30-62 does not anticipate or obviate the claimed step and means for comparing the expiration data of each of a plurality of actual computing systems to at least one target planning date for servicing the actual computing systems. What Yang mentions in the cited passage is response times for modeled systems. Again, since Yang is operating in the theoretical environment, real factors such as servicing dates are simply not relevant. Yang does not teach or suggest any consideration of target planning dates for servicing systems. Clearly, therefore, Yang does not teach or suggest the language of Claims 6-9 and 15-18.

In rejecting the language of Claims 7 and 16, the Examiner cites Col. 25, lines 13-20 of Yang. The cited passage states that "the transaction mix can be altered by changing the setting for number of transactions, the frequency, or the percentage mix for each transaction type." Those teachings relate to Yang generating a set of CAE for (i.e., modeling) a transaction. There is nothing in the cited passage which teaches or suggests expiration dates, target service dates for system servicing, or altering current workloads for at least two actual computing systems.

With regard to the language of Claims 8 and 17, the Examiner again cites the Yang teachings found in Col. 6, lines 9-36. The cited passage states that workload definition information for use in modeling can be collected from actual workload information, that the workload definition is "transferable to other applications and hardware", that "the same workload definition can be used for analysis performed with respect to a plurality of different hardware platforms". Creating a generic workload definition for system modeling does not anticipate the claims which recite creating "From" and "To" lists relative to at least one target service date for servicing of actual computing systems, and reassigning actual workload units based on assigned impact values reflecting a change in system expiration date if a workload unit is moved from one

YOR920000461US1 -32-

actual computing system to another. Yang is providing a generic workload definition that can be used to evaluate different candidate system. Such is not the same as or suggestive of evaluating and moving actual workload units in a runtime environment.

The language of Claims 9 and 18, which recites calculating new expiration dates for computing systems on the "From" and "To" lists after reassignment of workload units, has additionally been rejected. The Examiner cites Col. 5, line 1 through Col. 6, line 36 of Yang. The cited teachings all relate to how Yang defines its theoretical workload and CAE. Applicants again argue that the Yang teachings do not anticipate system expiration dates at all, since it is **modeling** candidate systems, and do not anticipate calculating new expiration dates of actual systems as claimed. Applicants reiterate that Yang is essentially designing a system and opts to change the system design, but does not teach or suggest altering the current workload of an existing system.

With regard to Claim 11, Applicants acknowledge that Fig. 7 of Yang illustrates a storage location. However, Yang's provision of a storage location in a "mechanism for profiling a system" is not the same as or suggestive of providing a storage location in the apparatus of Claim 10, including an administrative processor comprising an impact

YOR920000461US1 -33-

value component for assigning a plurality of impact values, one impact value for each workload unit assigned to each of the plurality of computing systems in a processing environment having a plurality of computer systems each having a plurality of assigned workload units, wherein said assigning of each impact value comprises determining a change in system expiration date should a workload unit be removed from the system; and a processing component for assessing the workload based on said impact values.

Appellants reiterate that the prior art must teach or suggest all of the claim limitations. "All words in a claim must be considered in judging the patentability of that claim against the prior art" (In re Wilson, 424 F. 2d 1382, 1385, 165 U.S.P.Q. 494, 496 (C.C.P.A. 1970)). Since the cited Yang and Chen references fail to teach each and every one of the claim limitations, a ***prima facie*** case of obviousness has not been established by the Examiner. Further, obviousness cannot be maintained without some teaching or suggestion of the claim features. The Federal Circuit has stated that when patentability turns on the question of obviousness, the obviousness determination "must be based on objective evidence of record" and that "this precedent has been reinforced in myriad decisions, and cannot be dispensed with." (In re Lee, 277 F. 3d 1338, 1343 (Fed. Cir. 2002)). Moreover, the Federal Circuit has stated

YOR920000461US1 -34-

that "conclusory statements" by an examiner fail to adequately address the factual question of motivation, which is material to patentability and cannot be resolved "on subjective belief and unknown authority" (Id. at 1343-1344). Applicants respectfully assert that the Examiner has not cited actual teachings but has provided conclusory statements about the teachings of Yang, Chen and Miller without basis in actual teachings from those patents.

Based on the foregoing amendments and remarks, Applicants respectfully request entry of the amendments, reconsideration of the amended claim language in light of the remarks, withdrawal of the rejections, and allowance of the claims.

Respectfully submitted,
L. B. Mummert, et al

/Anne Vachon Dougherty/
By: Anne Vachon Dougherty
Registration No. 30,374
Tel. (914) 962-5910